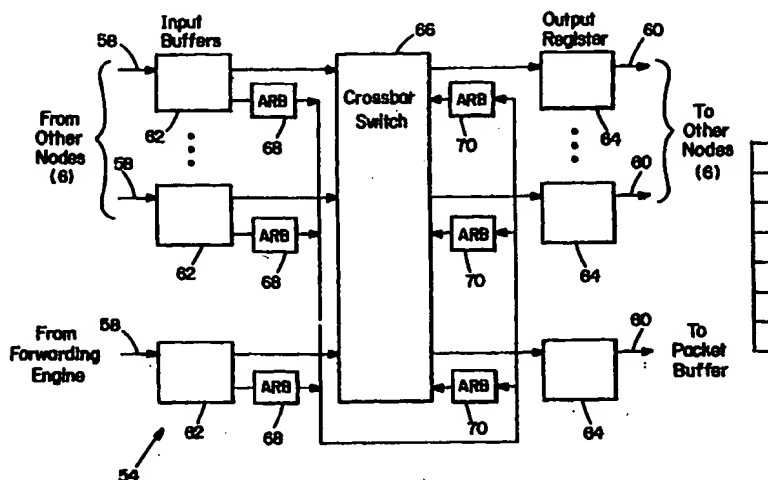




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04L 12/56		A1	(11) International Publication Number: WO 99/11033
			(43) International Publication Date: 4 March 1999 (04.03.99)
(21) International Application Number: PCT/US98/16762		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, HR, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 20 August 1998 (20.08.98)			
(30) Priority Data: 08/918,556 22 August 1997 (22.08.97) US 09/084,636 26 May 1998 (26.05.98) US			
(71) Applicant: AVICI SYSTEMS [US/US]; 101 Billerica Avenue, Billerica, MA 01826 (US).			
(72) Inventors: DALLY, William, J.; 1068 Vernier Place, Stanford, CA 94305 (US). CARVEY, Philip, P.; 7 Daniels Drive, Bedford, MA 01730 (US). DENNISON, Larry, R.; 505 Nahatan Street, Norwood, MA 02062 (US). KING, P., Allen; 30 Gibson Street, Needham, MA 02192 (US).			
(74) Agents: SMITH, James, M. et al.; Hamilton, Brook, Smith & Reynolds, P.C., Two Militia Drive, Lexington, MA 02421 (US).		<p>Published</p> <p><i>With international search report.</i></p> <p><i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>	

(54) Title: ROUTER WITH VIRTUAL CHANNEL ALLOCATION



(57) Abstract

An internet router is implemented as a network fabric of fabric routers and links. The internet router receives data packets from trunk lines or other internet links and analyzes header information in the data packets to route the data packets to output internet links. The line interface also analyzes the header to define a fabric path through the router fabric. The internet packets are broken into flits which are transferred through the router according to a wormhole routing protocol. Flits are stored in fabric routers at storage locations assigned to virtual channels corresponding the destination internet links. The virtual channels and links within the fabric define virtual networks in which congestion in one virtual network is substantially nonblocking to data flow through other virtual networks. Arbitration is performed at each fabric router to assign packets to virtual channels and to assign virtual channels to output fabric links. Data packets share access to multiple virtual channels, access being defined by dispersion code vectors. A virtual channel is enabled for possible assignment to an output fabric link upon receipt of an indication that an input buffer is available at the next fabric router of the path.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

ROUTER WITH VIRTUAL CHANNEL ALLOCATION

BACKGROUND OF THE INVENTION

Data communication between computer systems for applications such as web browsing, electronic mail, file transfer, and electronic commerce is often performed using a family of protocols known as IP (internet protocol) or sometimes TCP/IP. As applications that use extensive data communication become more popular, the traffic demands on the backbone IP network are increasing exponentially. It is expected that IP routers with several hundred ports operating with aggregate bandwidth of Terabits per second will be needed over the next few years to sustain growth in backbone demand.

As illustrated in Figure 1, the Internet is arranged as a hierarchy of networks. A typical end-user has a workstation 22 connected to a local-area network or LAN 24. To allow users on the LAN to access the rest of the internet, the LAN is connected via a router R to a regional network 26 that is maintained and operated by a Regional Network Provider or RNP. The connection is often made through an Internet Service Provider or ISP. To access other regions, the regional network connects to the backbone network 28 at a Network Access Point (NAP). The NAPs are usually located only in major cities.

The network is made up of links and routers. In the network backbone, the links are usually fiber optic communication channels operating using the SONET (synchronous optical network) protocol. SONET links operate at a variety of data rates ranging from OC-3 (155Mb/s) to OC-192 (9.9Gb/s). These links, sometimes called trunks, move data from one point to another, often over considerable distances.

Routers connect a group of links together and perform two functions: forwarding and routing. A data packet arriving on one link of a router is forwarded by sending it out on a different link depending on its eventual destination and the state of the output links. To compute the output link for a given packet, the router

participates in a routing protocol where all of the routers on the Internet exchange information about the connectivity of the network and compute routing tables based on this information.

Most prior art Internet routers are based on a common bus (Figure 2) or a crossbar switch (Figure 3). In the bus-based switch of Figure 2, for example, a given SONET link 30 is connected to a line-interface module 32. This module extracts the packets from the incoming SONET stream. For each incoming packet, the line interface reads the packet header, and using this information, determines the output port (or ports) to which the packet is to be forwarded. To forward the packet, the line interface module arbitrates for the common bus 34. When the bus is granted, the packet is transmitted over the bus to the output line interface module. The module subsequently transmits the packet on an outgoing SONET link 30 to the next hop on the route to its destination.

Bus-based routers have limited bandwidth and scalability. The central bus becomes a bottleneck through which all traffic must flow. A very fast bus, for example, operates a 128-bit wide datapath at 50MHz giving an aggregate bandwidth of 6.4Gb/s, far short of the Terabits per second needed by a backbone switch. Also, the fan-out limitations of the bus interfaces limit the number of ports on a bus-based switch to typically no more than 32.

The bandwidth limitation of a bus may be overcome by using a crossbar switch as illustrated in Figure 3. For N line interfaces 36, the switch contains $N(N-1)$ crosspoints, each denoted by a circle. Each line interface can select any of the other line interfaces as its input by connecting the two lines that meet at the appropriate crosspoint 38. To forward a packet with this organization, a line interface arbitrates for the required output line interface. When the request is granted, the appropriate crosspoint is closed and data is transmitted from the input module to the output module. Because the crossbar can simultaneously connect many inputs to many outputs, this organization provides many times the bandwidth of a bus-based switch.

Despite their increased bandwidth, crossbar-based routers still lack the scalability and bandwidth needed for an IP backbone router. The fan-out and fan-in required by the crossbar connection, where every input is connected to every output,

limits the number of ports to typically no more than 32. This limited scalability also results in limited bandwidth. For example, a state-of-the-art crossbar might operate 32 32-bit channels simultaneously at 200MHz giving a peak bandwidth of 200Gb/s. This is still short of the bandwidth demanded by a backbone IP router.

5 SUMMARY OF THE INVENTION

While they have limited bandwidth and scalability, crossbar-based routers have two desirable features:

1. They are non-blocking. As long as no two inputs request to communicate with the same output, all inputs can be simultaneously connected to their requested outputs. If one output becomes congested, the traffic to that output does not interfere with traffic addressed to other outputs.
2. They provide stiff backpressure. The direct connection between source and destination over the crossbar usually includes a reverse channel that may be used for immediate flow control. This backpressure can be used, for example, by an overloaded destination to signal a source to stop sending data.

To meet the requirements of routing for the internet backbone we would like to preserve these two properties while providing orders of magnitude greater bandwidth and scalability.

- 20 In accordance with one aspect of the present invention, advantages of crossbar-based internet routers are obtained with greater bandwidth and scalability by implementing the router itself as a multi-hop network.

An internet router embodying the invention receives data packets from a plurality of internet links and analyzes header information in the data packets to route the data packets to output internet links. The internet router comprises a fabric of fabric links joined by fabric routers, the number of fabric links to each fabric router being substantially less than the number of internet links served by the internet router. The fabric links and fabric routers provide data communication

between internet links through one or more hops through the fabric. In one embodiment, for example, 600 internet links are served by a 6 x 10 x 10 3-dimensional torus fabric array.

By providing a plurality of buffers in each fabric router, virtual channels
5 which share fabric output links may be defined. The virtual channels and links form a virtual network between internet router inputs and outputs in which congestion in one virtual network is substantially non-blocking to data flow through other virtual networks. A line interface to each internet link analyzes the header information in data packets received from the internet link to identify output internet links through
10 an internet routing protocol. The line interface further determines, through a fabric routing protocol, a routing path through the fabric to the identified output internet link. The packets may be subdivided into segments or flits (flow control digits) at the line interface, and those segments are forwarded through the fabric using wormhole routing. The line interface may define the routing path through the fabric
15 by including, in a header, a link definition of each successive link in the routing path. Each fabric router along the routing path stores an associated link definition from the header for forwarding successive segments of the packet.

Preferably, between hops on fabric links, flits are stored in fabric routers at storage locations assigned to virtual channels which correspond to destination
20 internet links. In one embodiment, the set of destination internet links is partitioned into disjoint subsets, and each virtual channel is assigned exclusively to one subset of destination internet links. In preferred embodiments, the number of internet links served by the internet router is at least an order of magnitude greater than the number of fabric links to each fabric router, and the number of virtual channels per
25 fabric router is substantially greater than the number of links to the fabric router.

To share virtual channels among data packets and to share fabric links among virtual channels, an arbitration is performed at each fabric router to assign a packet to a virtual channel for output from the fabric router and to assign a virtual
channel to an output fabric link from the fabric router. For flow control, a virtual
30 channel is enabled for possible assignment to an output fabric link upon receipt of an indication that an input buffer is available at the opposite end of the link.

Another aspect of the present invention which is applicable to internet routers and other network routers, including multicomputers, relates to an event-driven technique for handling virtual channels in a router that routes data packets. The router includes input physical channels that receive portions of the data packets, 5 output physical channels, and data buffers that are coupled with the input and output physical channels. The data buffers store the portions of the data packets. The router further includes control circuitry that is coupled with the input and output physical channels and the data buffers. The control circuitry generates channel assignments in response to queued events, and outputs the portions of the data 10 packets through the output physical channels according to the generated channel assignments. Preferably, the control circuitry assigns virtual channels to the data packets, assigns the output physical channels to the virtual channels in response to the queued events. In one embodiment, the router further includes a line interface coupled with an input physical channel and an output physical channel such that the 15 router forms an internet switch fabric router. In another embodiment, the router further includes a multicomputer interface coupled with an input physical channel and an output physical channel such that the router forms a multicomputer router for a multicomputer system.

According to the preferred embodiment, the control circuitry includes output 20 controllers that correspond to the output physical channels. Each output controller has a state table that records states of output virtual channels, and identifies input virtual channels connected with the output virtual channels. The input virtual channels hold the portions of the data packets.

Each output controller further includes an arbiter that is adapted to select 25 arrival events from multiple arrival queues, and state table logic that accesses that output controller's state table to assign output virtual channels in response to the selected arrival events. Each state table includes state vectors that correspond to output virtual channels.

Each state vector includes a busy indication that indicates whether that state 30 vector's corresponding output virtual channel is assigned to a data packet. Additionally, each state vector includes a wait field that indicates which of the input physical channels have received at least portions of data packets awaiting

assignment to that state vector's corresponding output virtual channel. Each wait field further indicates an order in which the input physical channels received the data packet portions. Each state vector further includes a present field that indicates a number of portions of a data packet present for transferring through that state
5 vector's output virtual channel to a downstream router. Furthermore, each state vector includes a credit field that indicates an amount of buffer space available in a downstream router coupled to that state vector's corresponding output virtual channel.

Each output controller further includes a transport circuit that queues
10 transport requests when that output controller's state table is accessed in response to the queued events, and forwards data packets through that output controller's output physical channel according to the queued transport requests. The portions of the data packets are flits of the data packets. Each transport circuit transmits a flit in response to a queued transport request.

15 Each output controller receives credit events from a downstream router and, in response to the received credit events, queues a transport request to transport a portion of a data packet over the corresponding output physical channel. In one embodiment, the queued events include tail credit events, and the output controllers free virtual channels only in response to the tail credit events.

20 The control circuitry can be shared by multiple virtual channels and activated to handle a particular virtual channel in response to an event.

Preferably, the control circuitry is adapted to generate virtual channel assignments that assign virtual channels to the data packets, and generate physical channel assignments that assign the output physical channels to the virtual channels.
25 Each of the assignments can be generated in response to queued arrival and credit events. The portions of the data packets are forwarded from the data buffers to the output physical channels according to the generated virtual and physical channel assignments.

In the primary implementation, to assure continued flow of data packets
30 despite congestion to some destinations, each destination node is assigned a virtual network comprising a unique virtual channel in each router. Any data packet destined to a particular output is stored in a data buffer corresponding to a unique

-7-

virtual channel for that destination, so congestion which blocks one data packet does not necessarily block another.

That approach is costly in buffer space and has limited scalability. In accordance with another aspect of the present invention, which may be applied to
5 any router without limitation to fabric routers, virtual networks are constructed with overlapping buffer assignments by sharing virtual channels. However, to prevent congestion in one virtual network from blocking transmissions in another virtual network, virtual networks share less than all virtual channels.

In accordance with this aspect of the present invention, a router for routing
10 data packets comprises input physical channels, output physical channels and data buffers, coupled with the input and output physical channels, for storing at least portions of the data packets. Control circuitry coupled to the input and output physical channels and the data buffers generates assignments of data packets to virtual channels which are multiplexed onto physical channels. A first data packet
15 has access to a first set of virtual channels of size C_1 , and a second data packet has access to a second set of virtual channels of size C_2 . There is an intersection of the first and second sets of size S where $0 < S < C_1$ and $S < C_2$. As a result, data packets may share access to multiple virtual channels to which other data packets have access, a first data packet sharing some but not all virtual channels with a second
20 data packet.

Preferably, substantially all of the data packets route over a plurality of virtual networks, each virtual network i using a subset of size C_i of the virtual channels. Virtual channel subsets of plural virtual networks overlap, sharing no more than $S < C_i$ virtual channels.

25 In a preferred router, the first packet routes on a first virtual network and the second packet routes on a second virtual network. All packets routed on the first virtual network share the first set of virtual channels and all packets routed on the second virtual network share the second set of virtual channels.

Preferably, each virtual network carries packets destined for a particular set
30 of destination nodes such as a pair of nodes opposite one another in a network. At least one of the virtual channels in each set of virtual channels can be determined by

a destination node address, specifically by a dimensional coordinate of the destination node.

To prevent interdimension deadlock, specific turns may be disallowed in each virtual network. In one embodiment, two virtual networks are assigned to each set of destination nodes and different turns are disallowed in each of the two virtual networks. In another embodiment, a single virtual network is assigned to each set of destination nodes. Multiple turns may be disallowed in that single virtual network.

In a preferred embodiment, the control circuitry includes output controllers that correspond to the output physical channels. Each output controller has a state table that records states of output virtual channels and identifies input virtual channels holding the portions of the data packets. State table logic accesses the state table to assign output virtual channels. The state table may include a virtual channel allocation vector for each virtual network and a busy vector which indicates virtual channels which are busy in the output physical channel. An output virtual channel is selected from a combination of the virtual channel allocation vector and the busy vector.

One application of the router is as a router within a multicomputer network, and another application is as a fabric router within an internet packet router.

20 BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, features and advantages of the invention will be apparent from the following more particular description of preferred embodiments of the invention, as illustrated in the accompanying drawings in which like reference characters refer to the same parts throughout the different views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention.

Figure 1 illustrates an internet configuration of routers to which the present invention may be applied.

Figure 2 is a prior art bus-based internet router.

30 Figure 3 is a prior art crossbar switch internet router.

Figure 4 illustrates a two-dimensional torus array previously used in direct multiprocessor networks.

Figure 5 illustrates an indirect network.

Figure 6 illustrates tree saturation of a network.

5 Figure 7 illustrates a three-dimensional fabric embodying the present invention.

Figure 8 illustrates the line interface module of a node in the array of Figure 7.

10 8. Figure 9 illustrates a fabric router used in the embodiment of Figures 7 and 8.

Figures 10A and 10B illustrate buffers, registers and control vectors used in the router of Figure 9.

Figures 11A and 11B illustrate alternative allocation control logic provided in input and output controllers, respectively, of the router of Figure 9.

15 Figure 12 illustrates a virtual channel state table used in the router of Figure 9.

Figure 13 illustrates a loop used to demonstrate dispersion routing.

Figure 14 illustrates turns required to reach a destination from each quadrant about the destination.

20 Figure 15 illustrates a virtual channel selection circuit embodying the invention.

DETAILED DESCRIPTION OF THE INVENTION

In implementing an internet router, the present invention borrows from multiprocessor technology and modifies that technology to meet the unique
25 characteristics and requirements of internet routers. In particular, each internet router is itself configured as either a direct or indirect network.

Multicomputers and multiprocessors have for many years used direct and indirect interconnection networks to send addresses and data for memory accesses between processors and memory banks or to send messages between processors.
30 Early multicomputers were constructed using the bus and crossbar interconnects shown in Figures 2 and 3. However, to permit these machines to scale to larger

numbers of processors they switched to the use of direct and indirect interconnection networks.

A direct network, as illustrated in Figure 4, is comprised of a set of processing nodes 40, each of which includes a router, R, along with a processor, P,
5 and some memory, M. These multicomputer routers should not be confused with the IP routers described above. They perform only forwarding functions and only in the very constrained environment of a multicomputer interconnection network. Each multicomputer router has some number, four in the example, of connections to other routers in the network. A processing node may send a message or make a
10 memory access to any other node in the system. It is not limited to communicating only with the immediately adjacent nodes. Messages to nodes that are further away are forwarded by the routers along the path between the source and destination nodes.

The network shown in Figure 4 is said to be direct since the channels are
15 made directly between the processing nodes of the system. In contrast, Figure 5 shows an indirect network in which the connections between process nodes 42 are made indirectly, via a set of router-only switch nodes 44. Direct networks are generally preferred for large machines because of the scalability. While an indirect network is usually built for a fixed number of nodes, a direct network grows with the
20 nodes. As more nodes are added, more network is added as well since a small piece of the network, one router, is included within each node.

Multicomputer networks are described in detail in Dally, W.J., "Network and Processor Architectures for Message-Driven Computing," VLSI and PARALLEL COMPUTATION, Edited by Suaya and Birtwistle, Morgan Kaufmann Publishers,
25 Inc., 1990, pp. 140-218. It should be stressed that multicomputer networks are local to a single cabinet or room as opposed to the Internet backbone network which spans the continent.

Direct and indirect multicomputer networks are scalable. For most common topologies the fan-in and fan-out of each node is constant, independent of the size of
30 the machine. Also, the traffic load on each link is either constant or a very slowly increasing function of machine size. Because of this scalability, these networks

-11-

have been successfully used to construct parallel computers with thousands of processing nodes.

Unfortunately, while multicomputer networks are scalable, they give up the two properties of crossbar networks that were crucial to IP switching: non-blocking behavior and stiff backpressure. Most economical direct and indirect networks are blocking. Because links are shared between multiple source-destination pairs, a busy connection between a pair of nodes can block the establishment of a new connection between a completely separate pair of nodes. Because packets in multicomputer networks are forwarded over multiple links with considerable queuing at each link, the backpressure, if any, from an overloaded destination node to a transmitting source node is late and soft if present at all.

The blocking nature of these switches and the soft nature of this backpressure is not a problem for a multicomputer because multicomputer traffic is self-throttling. After a processor has sent a small number of messages or memory requests (typically 1-8), it cannot send any further messages until it receives one or more replies. Thus, when the network slows down because of blocking or congestion, the traffic offered to the network is automatically reduced as the processors stall awaiting replies.

An IP switch, on the other hand, is not self-throttling. If some channels in the network become blocked or congested, the offered traffic is not reduced. Packets continue to arrive over the input links to the switch regardless of the state of the network. Because of this, an IP switch or router built from an unmodified multicomputer network is likely to become tree-saturated, and deny service to many nodes not involved in the original blockage. Moreover transient conditions often exist in IP routers where, due to an error in computing routing tables, a single output node can be overloaded for a sustained period of time. This causes no problems with a crossbar router as other nodes are unaffected. With a multicomputer network, however, this causes tree saturation.

Consider the situation illustrated in Figure 6. A single node in a 2-dimensional mesh network, node (3,3) labeled a, is overloaded with arriving messages. As it is unable to accept messages off the channels at the rate they are arriving, all four input channels to the node (b,a), (c,a), (d,a), (e,a), become

-12-

congested and are blocked. Traffic arriving at nodes b-e that must be forwarded across these blocked links cannot make progress and will back up along the edges into nodes b-e. For example, traffic into node b backs up along (f,b), (g,b), and (h,b). If the blockage persists, the channels into f-h and related nodes become
5 blocked as well and so on. If the overload on node a persists, eventually most of the channels in the network will become blocked as a tree of saturation expands outward from node a.

The major problem with tree saturation is that it affects traffic that is not destined for node a. A packet from (1,4) to (5,3) for example may be routed along a
10 path (dotted line) that includes (f,b) and (b,a) for example. Since these links are blocked, traffic from node (1,4) to node (5,3) is blocked even though neither of these nodes is overloaded.

The router of the present invention overcomes the bandwidth and scalability limitations of prior-art bus- and crossbar- based routers by using a multi-hop
15 interconnection network, in particular a 3-dimensional torus network, as a router. With this arrangement, each router in the wide-area backbone network in effect contains a small in-cabinet network. To avoid confusion we will refer to the small network internal to each router as the switching fabric and the routers and links within this network as the fabric routers and fabric links.

20 Unlike multicomputer networks, the switching fabric network is non-blocking and provides stiff backpressure. These crossbar-like attributes are achieved by providing a separate virtual network for each destination node in the network.

Typical packets forwarded through the internet range from 50 bytes to 1.5 Kbytes. For transfer through the fabric network of the internet router of the present
25 invention, the packets are divided into segments, or flits, each of 36 bytes. At least the header included in the first flit of a packet is modified for control of data transfer through the fabric of the router. In the preferred router, the data is transferred through the fabric in accordance with a wormhole routing protocol.

Each virtual network comprises a set of buffers. One or more buffers for
30 each virtual network are provided on each node in the fabric. Each buffer is sized to hold at least one flow-control digit or flit of a message. The virtual networks all share the single set of physical channels between the nodes of the real fabric

-13-

network. A fair arbitration policy is used to multiplex the use of the physical channels over the competing virtual networks. Each virtual network has a different set of buffers available for holding the flits of its messages.

For each pair of virtual networks A and B, the set of buffers assigned to A
5 contains at least one buffer that is not assigned to B. Thus if network B is blocked, A is able to make progress by forwarding messages using this buffer that is not shared with B although it may be shared with some other virtual network.

One simple method for constructing virtual networks is to provide a separate flit buffer, a virtual channel, on each node for each virtual network and thus for each
10 destination. For example, in a machine with $N=512$ nodes and hence 512 destinations, each node would contain 512 distinct flit buffers. Buffer i on each node is used only to hold flits of messages destined for node i . This assignment clearly satisfies the constraints above as each virtual network is associated with a singleton set of buffers on each node with no sharing of any buffers between virtual
15 networks. If a single virtual network becomes congested, only its buffers are affected, and traffic continues on the other virtual networks without interference. An alternative dispersive approach is discussed below.

The preferred router is a 3-dimensional torus network of nodes as illustrated in Figure 7. Each node N comprises a line interface module that connects to
20 incoming and outgoing SONET internet links. Each of these line-interface nodes contains a switch-fabric router that includes fabric links to its six neighboring nodes in the torus. IP packets that arrive over one SONET link, say on node A, are examined to determine the SONET link on which they should leave the internet router, say node B, and are then forwarded from A to B via the 3-D torus switch
25 fabric.

The organization of each node or line-interface module is illustrated in Figure 8. Packets arrive over the incoming SONET link 46, and the line interface circuit 48 converts the optical input to electrical signals and extracts the packets and their headers from the incoming stream. Arriving packets are then passed to the
30 forwarding engine hardware 50 and are stored in the packet memory 52. The forwarding engine uses the header of each packet to look up the required output link for that packet. In conventional IP router fashion, this lookup is performed by

-14-

traversing a tree indexed by the header fields. The leaves of the tree contain the required output link, as in a conventional IP router, and additionally include the route through the switch fabric to the output link. Finally, the packet along with its destination and route are passed to the fabric router 54 of the node for forwarding
5 through the fabric to the output node. From the fabric router 54 of the output node, the packet is delivered through the packet buffer 52 of that node and through the line interface circuit 48 to the output link 56.

Packets in the internet router are forwarded from the line-interface module associated with the input trunk to the line-interface module associated with the
10 output trunk using source routing. With source routing, the route of links through intermediate fabric routers is determined by a table lookup in the input module. This lookup is performed by the forwarding engine before presenting the packet to the fabric router. Alternative paths allow for fault tolerance and load balancing.

The source route is a 10-element vector where each element is a 3-bit hop
15 field. Each hop field encodes the output link to be taken by the packet for one step of its route, one of the six inter-node links or the seventh link to the packet buffer of the present node. The eighth encoding is unused. This 10-element vector can be used to encode all routes of up to 10 hops which is sufficient to route between any pair of nodes in a 6 x 10 x 10 torus. Note that all 10 elements need not be used for
20 shorter routes. The last used element selects the link to the packet buffer 52 or may be implied for a 10-hop route.

As the packet arrives at each fabric node along the route, the local forwarding vector entry for that packet is set equal to the leftmost element of the source route. The source route is then shifted left three bits to discard this element
25 and to present the next element of the route to the next router. During this shift, the 3-bit code corresponding to the packet buffer of the present node is shifted in from the right. Subsequent flits in that packet follow the routing stored for that packet in the router.

One skilled in the art will understand that there are many possible encodings
30 of the fabric route. In an alternative embodiment, the fact that packets tend to travel in a preferred direction in each dimension may be exploited to give a more compact encoding of the fabric

route. In this embodiment, the route is encoded as a three-bit preferred direction followed by a multiplicity of two-bit hop fields. The three-bit field encodes the preferred direction (either positive or negative) for each dimension of the network (x, y, and z). For each step or hop of the route, a two-bit field selects the dimension
5 over which the next hop is to be taken (0=x, 1=y, or 2=z). The direction of this hop is determined by the preferred direction field. The fourth encoding of the two-bit hop field (3) is used as an escape code. When a hop field contains an escape code, the next hop field is used to determine the route. If this second hop field contains a dimension specifier (0-2), the hop is taken in the specified dimension in the direction
10 opposite to the preferred direction and the preferred direction is reversed. If the second hop field contains a second escape code, the packet is forwarded to the exit port of the fabric router. With this encoding, as packets arrive at a fabric node, the local forwarding vector entry for that packet is computed from the preferred direction field and the leftmost hop field. The hop fields are then shifted left two
15 bits to discard this field and to present the next field to the next router. During this shift, the two-bit escape code is shifted into the rightmost hop field. For packets that travel primarily in the preferred direction, this encoding results in a more compact fabric route as only two bits, rather than three, are needed to encode each hop of the route.

20 A fabric router used to forward a packet over the switch fabric from the module associated with its input link to the module associated with its output link is illustrated in Figure 9. The router has seven input links 58 and seven output links 60. Six of the links connect to adjacent nodes in the 3-D torus network of Figure 7. The seventh input link accepts packets from the forwarding engine 50 and the
25 seventh output link sends packets to the packet output buffer 52 in this router's line interface module. Each input link 58 is associated with an input buffer 62 and each output link 60 is associated with an output register 64. The input buffers and output registers are connected together by a 7 x 7 crossbar switch 66.

One skilled in the art will understand that the present invention can be
30 practiced in fabric networks with different topologies and different numbers of dimensions. Also, more than one link may be provided to and from the line interface. In an alternative embodiment two output links are provided from the

-16-

5 fabric to the line interface bringing the total number of output links, and hence output registers, to eight. In this case, the input buffers and output registers are connected by a 7x8 crossbar switch. The second output link provides additional bandwidth to drain packets from the fabric network when a single node receives traffic simultaneously from many directions.

10 A virtual network is provided for each pair of output nodes. Each of the seven input buffers 62 contains a buffer, of for example one flit, for each virtual network in the machine. In one embodiment, a 6 x 10 x 10 torus fabric provides 600 nodes. A single virtual network is assigned to a pair of maximally distant output nodes in the network as minimal routes between these two nodes are guaranteed not to share any links and thus are guaranteed not to interfere with one another. Further, two virtual networks are provided for each pair of nodes to allow for two priorities in serving different classes of traffic. Thus, in the router, there are 600 virtual networks: two virtual networks for each of 300 pairs of nodes. Each input buffer 62
15 contains space for 600 36-byte flits (21,600 bytes total).

As an improvement, each input buffer has storage for two flits for each virtual channel. The size of a flit determines the maximum duty factor of a single virtual channel and the fragmentation loss associated with rounding up packets to a whole number of flits. The maximum bandwidth on a single fabric link that can be
20 used by a single virtual channel can be no more than the flit size times the number of flits per virtual channel buffer divided by the time for a header flit to propagate through a router. For example, if a flit is 36 Bytes, there is a single flit per buffer, and it takes ten 10ns clocks for a header flit to propagate through a router, the maximum bandwidth per virtual channel is 360MBytes/s. If the link bandwidth is
25 1200MBytes/s, a single virtual channel can use at most 30% of the link bandwidth. If the flit buffer capacity is at least as large as the link bandwidth divided by the router latency (120Bytes in this case), a single virtual channel can use all of the link capacity.

30 One would like to make the flit size as large as possible both to maximize the link bandwidth that a single virtual channel can use and also to amortize the overhead of flit processing over a larger payload. On the other hand, a large flit reduces efficiency by causing internal fragmentation when small packets must be

rounded up to a multiple of the flit size. For example, if the flit size is 64 Bytes, a 65 Byte packet must be rounded up to 128 Bytes, incurring nearly 50% fragmentation overhead.

One method for gaining the advantages of a large flit size without incurring the fragmentation overhead is to group adjacent flits into pairs that are handled as if they were a single double-sized flit. For all but the last flit of an odd-length message, all flit processing is done once for each flit pair, halving the flit processing overhead. The last odd flit is handled by itself. However, these odd single-flits are rare and so their increased processing overhead is averaged out. In effect, flit pairing is equivalent to having two sizes of flits – regular sized and double sized. The result is that long messages see the low processing overhead of double-sized flits and short messages see the low fragmentation overhead of regular sized flits. In the preferred embodiment, flits are 36 Bytes in length and are grouped into pairs of 72 Bytes total length.

If a virtual channel of a fabric router destined for an output node is free when the head flit of a packet arrives for that virtual channel, the channel is assigned to that packet for the duration of the packet, that is, until the worm passes. However, multiple packets may be received at a router for the same virtual channel through multiple inputs. If a virtual channel is already assigned, the new head flit must wait in its flit buffer. If the channel is not assigned, but two head flits for that channel arrive together, a fair arbitration must take place. With limited buffer space assigned to each virtual channel, a block at an output node from the fabric is promptly seen through backpressure to the input line interface for each packet on that virtual network. The input line interface can then take appropriate action to reroute subsequent packets. With assignment of different destinations to different virtual networks, interference between destinations is avoided. Traffic is isolated.

Once assigned an output virtual channel, a flit is not enabled for transfer across a link until a signal is received from the downstream node that an input buffer at that node is available for the virtual channel.

An elementary flow control process is illustrated in Figures 9, 10A and 10B. Each cycle, a number M of the enabled flits in each input buffer are selected by a fair arbitration process 68 to compete for access to their requested output links. The

-18-

selected flits forward their output link requests to a second arbiter 70 associated with the requested output link. This arbiter selects at most one flit to be forwarded to each output link. The winning flits are then forwarded over the crossbar switch to the output register and then transmitted over the output link to the next router in the switch fabric. Until selected in this two-step arbitration process, flits remain in the input buffer, backpressure being applied upstream.

The fabric router at each line-interface module uses credit-based flow-control to regulate the flow of flits through the fabric network. Associated with each set of input buffers 62 are two V-bit vectors; a presence vector, P, and an enabled vector, E. V, as illustrated in Figure 10A, is the number of virtual networks and hence the number of entries in the buffer. A bit of the presence vector, $P[v,i]$, is set if the input buffer i contains a flit from virtual network v. Bit $E[v,i]$ is set if this flit is enabled to take the next hop of the route to its destination link.

As illustrated in Figure 10B, associated with each output register is a V-bit credit vector, C, that mirrors the complement of the presence vector on the opposite end of the fabric link at the receiving node. That is, $C[v,j]$ is set at a given output j if $P[v,i]$ is clear at the input port on the opposite side of the link. If $C[v,j]$ is set, then the output register has a credit for the empty buffer at the opposite end of the link.

Flits in an input buffer are enabled to take their next hop when their requested output link has a credit for their virtual network. For example, suppose the packet in virtual network v of input buffer i has selected output link j for the next hop of its route. We denote this as $F[v,i] = j$, where F is the forwarding vector. The flit in this input buffer is enabled to take its next hop when two conditions are met. First, it must be present, $P[v,i] = 1$, and second, there must be a credit for buffer space at the next hop, $C[v,j] = 1$.

The input buffer storage is allocated separately to each virtual network while the output registers and associated physical channels are shared by the virtual networks. The credit-based flow control method guarantees that a virtual network that is blocked or congested will not indefinitely tie up the physical channels since only enabled flits can compete in the arbitration for output links. Further, because only one or two flits per virtual network are stored in each input buffer, stiff

backpressure is applied from any blocked output node to the forwarding engine of the input node.

ALLOCATION

Arbitration and flow control can be seen as an allocation problem which
5 involves assigning virtual channels to packets, arriving from different input nodes and destined to common output nodes, and assigning physical channel bandwidth to flits destined to the same next node in the fabric path.

In a multistage switching fabric, packets composed of one or more flits advance from their source to their destination through one or more fabric routers. At
10 each hop, the head flit of a message arrives at a node on an input virtual channel. It can advance no further until it is assigned an output virtual channel. In the switch fabric of the preferred embodiment each packet may route on only one virtual channel. If the virtual channel is free when the packet arrives, it is assigned to the arriving packet. If, however, the virtual channel is occupied when the packet
15 arrives, the packet must wait until the output virtual channel becomes free. If one or more packets are waiting on a virtual channel when it is released, an arbitration is performed and the channel is assigned to one of the waiting packets.

Once a packet succeeds in acquiring the virtual channel it must compete for physical channel bandwidth to advance its flits to the next node of its route. A
20 packet can only compete for bandwidth when two conditions hold. First, at least one flit of the packet must be present in the node. Second, there must be at least one flit of buffer space available on the next node. If these two conditions do not hold, there is either no flit to forward or no space in which to put the flit at the next hop. If both conditions hold for a given packet, then that packet is enabled to transmit a flit.
25 However, before a flit can be sent, the packet must win two arbitrations. Among all the enabled packets, for a flit of the packet to advance to the next node of the route, a packet must be granted both an output port from the input flit buffer and the output physical channel.

For small numbers of virtual channels, the allocation problem can be solved
30 in parallel for the elementary case of figures 9, 10A and 10B using combinational logic.

-20-

Consider first the virtual channel allocation problem. A bit of state, H , is associated with each of V input virtual channels on each of K input controllers. This bit is set if the input virtual channel contains a head flit that has not yet been assigned an output virtual channel. The bit array $H[1:V, 1:K]$ determines the demand for virtual channels. A bit of state, B , is associated with each of V output virtual channels in each of K output controllers. This bit is set if the output virtual channel is busy. The bit array $B[1:V, 1:K]$ determines the allocation status of the virtual channels.

To allocate a virtual channel, v , in output controller, k , an arbitration must first be performed across virtual channel v in each of the k input controllers with input controller i only competing if (1) $H[v, i]$ is set and (2) the destination of the channel, $F[v, i] = k$. The input that wins the arbitration is granted the virtual channel only if $B[v, k] = 0$.

The situation is similar for allocation of physical channel bandwidth to flits. The buffer status of each input virtual channel is indicated by a presence bit, P , that is set when one or more flits are in the present node. Each output virtual channel looks ahead and keeps a credit bit, C , that is set when one or more empty buffers are available in the next node. Suppose we choose to do the allocation serially (which is sub-optimal); first arbitrating for an output port of the input controller and then arbitrating for an output channel. Suppose each input buffer has M output ports. Then for input buffer i , we first determine which virtual channels are enabled. An enabled vector, $E[v, i]$ is calculated as $E[v, i] = \neg H[v, i] \wedge P[v, i] \wedge C[v, j]$ where \neg denotes logical negation, \wedge denotes a logical AND operation, and j is the destination of the packet on virtual channel v of input controller i . Thus, a packet is enabled to forward a flit when it is not waiting for a virtual channel, when there is at least one flit present in its buffer, and when there is at least one flit of storage available at the next hop. Next, all of the enabled channels in the input buffer arbitrate for the M output ports of the input buffer. This requires a V -input M -output arbiter. Finally, the winners of each local arbitration arbitrate for the output virtual channels, this takes K , MK -input arbiters.

With large numbers of virtual channels a combinational realization of the allocation logic requires a prohibitive number of gates. The preferred switch fabric

-21-

has $V=600$ virtual channels and $K=7$ ports. To implement this allocation method having combinational logic thus requires 4200 elements of vectors H and B, 4200 3:8 decoders to qualify the arbitrations, and 4200 7-input arbiters to select the winners. Between the flip-flops to hold the state, the decoders, and the arbiters,
5 about 50 2-input gates are required for each of the 4200 virtual channels for a total of over 200,000 logic gates, a prohibitive number.

For the preferred router, the P and C arrays are also 4200 elements each. Between the C-multiplexers and the arbiters, each element requires about 40 gates. Thus the bandwidth allocation requires an additional 160,000 logic gates.

10 While quite reasonable for routers with small numbers of virtual channels, V less than or equal to 8, combinational allocation is clearly not feasible for the router with $V=600$.

EVENT-DRIVEN ALLOCATION

The logic required to perform allocation can be greatly reduced by observing
15 that for large numbers of virtual channels, the state of most virtual channels is unchanged from one cycle to the next. During a given flit interval, at most one virtual channel of a given input controller can have a flit arrive, and at most M virtual channels can have a flit depart. The remaining $V-M-1$ virtual channels are unchanged.

20 The sparse nature of changes to the virtual channel state can be exploited to advantage through the use of event-driven allocation logic. With this approach, a single copy (or a small number of copies) of the virtual channel state update, and allocation logic is multiplexed across a large number of virtual channels. Only active virtual channels, as identified by the occurrence of events, have their state
25 examined and updated and participate in arbitration.

Two types of events, arrival events and credit events, activate the virtual channel state update logic. A third type of event, a transport event, determines which virtual channels participate in arbitration for physical channel bandwidth. Each time a flit arrives at a node, an arrival event is queued to check the state of the
30 virtual channel associated with that flit. A similar check is made in response to a credit event which is enqueued each time the downstream buffer state of a virtual

-22-

channel is changed. Examining the state of a virtual channel may lead to allocation of the channel to a packet and/or scheduling a flit for transport to the downstream node. In the latter case, a transport event is generated and enqueued. Only virtual channels with pending transport events participate in the arbitration for input buffer
5 output ports and output physical channels. Once a flit wins both arbitrations and is in fact transported, the corresponding transport event is dequeued.

Logic to implement event-driven channel allocation is illustrated in Figures 11A and 11B. Figure 11A shows one of seven input controllers while Figure 11B shows one of seven output controllers. Each input controller is connected to each
10 output controller at the three points shown. Each input controller includes a destination table 72, an arrival queue 74, a credit queue 76 and a flit buffer 62. A virtual channel state table 80 and a transport queue 82 are included in each output controller. The Figures show an event-driven arrangement where the virtual channel state is associated with each output controller. It is also possible to associate the
15 state with the input controllers. Placing the state table in the output controller has the advantage that virtual channel allocation (which must be performed at the output controller) and bandwidth allocation (which can be performed at either end) can be performed using the same mechanism.

The destination tables, flit buffers, and virtual-channel state tables have
20 entries for each virtual channel, while the three queues require only a small number of entries. For each virtual channel, the destination table records the output port required by the current packet on that input channel, if any, (i.e., F_p), the flit buffer 62 provides storage for one or more flits of the packet, and the state of the output virtual channel is recorded in the state table. The arrival, credit, and transport
25 queues contain entries for each event that has occurred but has not yet been processed.

On the input side, the dual-ported arrival queue, credit queue, and flit buffer also serve as a synchronization point as illustrated by the dashed line in Figure 11A. The left port of these three structures, and all logic to the left of the dotted line
30 (including the destination table), operates in the clock domain of the input channel. The right port of these three structures, and all logic to the right of the dotted line, including Figure 11B, operate in the internal clock domain of the router.

In an alternative embodiment arriving flits are synchronized to the local clock domain before accessing the arrival queue or destination table.

With the arrangement shown in Figures 11A and 11B, an allocation of a virtual channel or a physical channel flit cycle is performed through a three-event sequence of arrival, transport, and credit. An arriving flit arbitrates for access to the state table for its output virtual channel. When granted, the table is updated to account for the arriving flit and, if the channel is allocated to its input controller and a credit is available, a transport request is queued to move the flit. The transport request arbitrates for access to the input flit buffer. When access is granted the flit is removed from the buffer and forwarded to the next node. Whenever a flit is removed from the flit buffer a credit is queued to be transmitted to the previous node. When credits arrive at a node, they update the virtual channel state table and enable any flits that are waiting on zero credits. Finally, the arrival of a tail flit at a node updates the virtual channel state to free the channel.

Each time a flit arrives at an input controller, the contents of the flit are stored in the flit buffer 62. At the same time, the destination table 72 is accessed, and an arrival event, tagged with the required output port number, is enqueued at 74. The destination table is updated by the head flit of each packet to record the packet's output port and then consulted by the remaining flits of a packet to retrieve the stored port number. An arrival event includes a virtual channel identifier (10 bits), a head bit, and an output port identifier (3 bits). The arrival events at the heads of each of the K input controller's arrival queues (along with input port identifiers (3 bits)) are distributed to arbiters 84 at each output controller. At each output controller the arrival events, that require that output port, arbitrate for access to the state table 80. Each cycle, the winning arrival events are dequeued and processed. The losing events remain queued and compete again for access to the state table on the subsequent cycle.

As shown in Fig. 12, for each output virtual channel, v, on output k, the virtual channel state table 80 maintains a state vector, $S[v,k]$ containing:

1. The allocation status of the channel, B, idle (0), busy (1) or tail pending (2).
2. The input controller assigned to this channel (if B is set), I, (3 bits).

-24-

3. A bit vector of input controllers waiting on this channel, W, (7 bits).
4. The number of credits (empty buffers on the next node), C, (1 bit).
5. The number of flits present on this node, P, (1 bit).

The first three of these (B,I,W) are associated with the allocation of output
 5 virtual channels to input virtual channels while the last two (C,P) are associated with
 the allocation of physical channel bandwidth to flits. The number of flits in each
 element of the state vector may be varied as appropriate. For example, if more flit
 buffers are available on each node, then more bits would be allocated to the C and P
 field. Much of the state here corresponds directly to the state bits in the
 10 combinational logic approach. The B, C, and P bits are identical. The W bits
 correspond to the H bits, qualified by required output channel.

The number of bits in the waiting vector, W, can be increased to provide
 improved fairness of arbitration. With just a single bit, a random or round-robin
 arbitration can be performed. If 3-bits are stored for each entry, a queuing
 15 arbitration can be performed with the input virtual channels serviced in the order that
 their requests arrived. Each virtual channel in effect "takes a number" when it
 arrives at the state table, and this number is stored in its entry of the W vector.
 When the channel becomes free, the "next" number is served.

When an arrival event associated with virtual channel v, from input
 20 controller I, arrives at the state table for output k, it reads S[v,k] and performs one of
 the following actions depending on the type of event (heads vs. body) and the state
 of the channel.

1. If the flit is a head, the channel is idle, B=0, and there are downstream
 credits, C≠0, (a) the channel is assigned to the input by setting B=1, I=i, (b)
 25 a downstream buffer is allocated by decrementing C, and (c) a transport
 request is queued for (v,i,k) at 82.
2. If the flit is a head, the channel is idle, but there are no downstream credits,
 the channel is assigned to the input, and the presence count, P, is
 incremented. No downstream buffer is allocated and no transport request is
 30 queued.

-25-

3. If the flit is a head and the channel is busy, $B=1$, the virtual channel request is queued by setting the i th bit of the wait vector, W .
4. If the flit is a body flit, and there are downstream credits, a downstream buffer is allocated and a transport request is queued.
5. If the flit is a body flit, and there are no downstream credits, the presence count is incremented.
6. If the flit is a tail and $W=0$, no waiting heads, then, if there is a credit available the tail flit is queued for transport and the channel is marked idle, $B=0$. Otherwise, if no credit is available, the channel is marked tail pending, $B=2$, so the arrival of a credit will transmit the tail and free the channel.
7. If the flit is a tail, a credit is available ($C \neq 0$), and there are packets waiting ($W \neq 0$), the tail flit is queued for transport as in cases 1 and 4 above. An arbitration is performed to select one of the waiting inputs, j . The channel is assigned to that input ($B=1$, $I=j$), and, if there is an additional credit available, this new head flit is queued for transport; otherwise it is marked present.
8. If the flit is a tail and a credit is not available, ($C=0$), the presence count is incremented and the status of the channel is marked "tail pending," ($B=2$).

If there is just a single flit buffer per virtual channel, when a body flit arrives there is no need to check the virtual channel allocation status (B , I and W) as the flit could only arrive if the channel were already allocated to its packet ($B=1$, $I=i$). If there is more than one flit buffer per virtual channel, the virtual channel of each body flit arrival must be checked. Flits arriving for channels that are waiting for an output virtual channel will generate events that must be ignored. Also, the number of flits buffered in a waiting virtual channel must be communicated to the state table when the output channel is allocated to the waiting channel. This can be accomplished, for example, by updating the flit count in the state table from the count in the flit buffer whenever a head flit is transported. Note that in case 1 above, we both allocate the virtual channel and allocate the channel bandwidth for the head flit in a single operation on the state table. Tail flits here result in a pair of actions: the tail flit is first processed as a body flit to allocate the bandwidth to move the tail

flit, the tail flit is then processed as a tail flit to free the channel and possibly move a pending head flit. Unless the transport queue can accept two inputs simultaneously, this must be done sequentially as a tail flit arrival may enqueue two flits for transport: the tail flit itself, and the head flit of a waiting packet.

- 5 Each entry in the transport queue (v,i,k) is a request to move the contents of flit buffer v on input controller i to output k . Before the request can be honored, it must first arbitrate at 86 for access to flit buffer i . On each cycle, the transport requests at the head of the queues in each of the K output controllers are presented to their requested input buffers where they arbitrate for access to the M ports. The
- 10 winning transport requests are dequeued and their flits forwarded to the appropriate output multiplexer 88. The other requests remain in the transport queues. There is no need to arbitrate for a fabric link here, as the output controller associated with each of the outgoing fabric links makes at most one request per cycle.

- Each time a transport request successfully forwards a flit to an output, a
- 15 credit is generated to reflect the space vacated in the input flit buffer. This credit is enqueued in a credit queue 76 for transmission to the output controller of the previous node. When a credit for virtual channel v arrives at output controller k of a node, it reads the state vector, $S[v,k]$, to check if any flits are waiting on credits. It proceeds as follows depending on the state of the presence count.

- 20 1. If there are no flits waiting, $P=0$, the credit count is incremented, $C=C+1$.
2. If there are flits waiting, $P \neq 0$, the number of waiting flits is decremented, $P=P-1$, and a transport request for the first waiting flit is enqueued.
3. If there is a tail flit pending ($B=2$), a transport request for the tail flit is
- 25 queued. If no head flits are waiting on the channel ($W=0$), the channel is set idle ($B=0$). Otherwise, if there are head flits waiting ($W \neq 0$), an arbitration is performed to select a waiting channel, say from input controller j , the channel is allocated to this channel ($B=1, I=j$), and the head flit is marked present ($P=1$) so the next arriving credit will cause the head flit to be transmitted.

-27-

In the above-described event-driven embodiment, the output controller processes body flits and tail flits differently. In particular, the output controller processes body flits according to techniques 4 and 5, and processes tail flits according to techniques 6, 7 and 8, described above.

5 As described in technique 7, a head flit of a data packet can follow directly on the heels of a tail flit of a previous data packet. For example, a data packet can occupy a virtual channel while one or more data packets (i.e., one or more head flits) wait for that virtual channel. When an arrival event for a tail flit of the occupying data packet reaches the output controller, the output controller queues the tail flit for
10 transmission to the next fabric router downstream, and allocates the virtual channel to one of the waiting data packets (i.e., one of the waiting head flits). Accordingly, the output controller grants the virtual channel to a new data packet as soon as the fabric router queues the tail flit for transmission.

In an alternative event-driven embodiment, the output controller processes
15 body flits and tail flits similarly. In particular, the output controller processes both body and tail flits according to techniques 4 and 5, as described above. As such, when an arrival event for a tail flit reaches the output controller, and when a credit is available, the output controller queues the tail flit for transmission without freeing the virtual channel or allocating the virtual channel to a waiting data packet. When a
20 fabric router that is downstream from the present fabric router receives, processes and forwards the tail flit, the downstream fabric router generates a special tail credit in place of the normal credit. The downstream fabric router sends this tail credit upstream to the present fabric router. When the output controller of the present fabric router receives the tail credit, the output controller increments the credit count
25 of the virtual channel in a manner similar to that for normal credits, and frees the virtual channel. At this point, if there are data packets waiting for the virtual channel, the output controller performs an arbitration procedure to assign the virtual channel to one of the waiting data packets.

The fabric router according to the alternative event-driven embodiment has
30 slower performance than the fabric router of the event-driven embodiment that processes body and tail flits differently. In particular, after the fabric router of the alternative embodiment queues a transport request for transmission of a tail flit to a

-28-

downstream router, the virtual channel assigned to the data packet of that tail flit becomes idle. The virtual channel is not available for use by another data packet until the fabric router receives a tail credit from the downstream fabric router.

However, the alternative event-driven embodiment results in considerably
5 simpler logic for several reasons. First, it simplifies the handling of events by reducing the complexity of handling a tail-flit arrival event. The work is instead spread between the tail-flit arrival and the tail-credit events. Furthermore, it simplifies the logic by ensuring that only a single packet is in a given virtual channel's flit buffer at any point in time. This is guaranteed by not granting the
10 virtual channel to a new packet until the tail of the previous packet has cleared the flit buffer - as signaled by the tail credit. In contrast, in the event-driven embodiment that processes body flits and tail flits differently, a head flit of a next packet can follow directly on the heels of the tail flit of a present packet, and two or more packets may be queued in a single virtual channel's flit buffer at the same time.

15 Each event-driven method of allocation described here reduces the size and complexity of the logic required for allocation in two ways. First, the state information for the virtual channels can be stored in a RAM array with over 10x the density of the flip-flop storage required by the combinational logic approach. Second, the selection and arbitration logic is reduced by a factor of V. Arbitration
20 for access to the virtual channel buffers is only performed on the channels for which changes have occurred (flit or credit arrival), rather than on all V channels.

Only the flit buffer, the state table, and the destination table in Figures 11A and 11B need to have V entries. A modest number of entries in the bid, transport, and credit queues will suffice to smooth out the speed mismatches between the
25 various components of the system. If a queue fills, operation of the unit filling the queue is simply suspended until an entry is removed from the queue. Deadlock can be avoided by breaking the cycle between event queues. For example, by dropping transport events when the transport queue fills, the state table is able to continue to consume credit and arrival events. Lost events can be regenerated by periodically
30 scanning the state table. Alternately, one of the N queues, e.g., the transport queue, can be made large enough to handle all possible simultaneous events, usually V times N (where N is the number of flits in each channel's input buffer).

DISPERSION

While assigning a separate virtual channel to each virtual network as described above is a simple solution, it is costly and has limited scalability. The number of buffers required in each interconnection network router increases linearly
5 with the number of nodes in the system. With 512 virtual networks, the number of flit buffers required is pushing the physical limits of what can be economically constructed on the integrated circuits making up the router's switch fabric. Also, with each virtual channel dedicated exclusively to a virtual network, if a virtual network is unused, the virtual channel remains unused, and underutilization of
10 virtual channels across the network results.

Further, it would be desirable to increase the number of flit buffers for each virtual channel at each node in order to increase the velocity of the network. In the design described above, two flit buffers are provided for each virtual channel in each node, but acknowledgments delay the transfer of flits between flit buffers. By
15 increasing the number of flit buffers per virtual channel per node, acknowledgments, can be made in parallel with transfers, and the velocity of the system can be increased.

To reduce the number of buffers, and hence the cost, of the switch fabric and to provide for greater scalability, velocity and utilization, virtual networks may be
20 constructed with overlapping buffer assignments by sharing virtual channels. In this way, the number of virtual channels required to serve all of the virtual networks can be substantially reduced. On the other hand, for the system to approach the behavior of a cross-bar switch, it is important that loss of a shared virtual channel due to congestion in one virtual network not block transmissions of another virtual
25 network. Accordingly, each virtual network must have access to multiple virtual channels at each node, and it is important that, for any two virtual networks x and y , the virtual network x have access to a virtual channel to which the virtual network y does not have access.

The shared virtual channel system can be implemented using dispersion
30 codes which disperse the assignment of virtual channels across the virtual networks. Consider for example a network with M nodes (and hence M virtual networks) which share N virtual channels on each node. Each virtual network, j , is assigned a

-30-

dispersion code, an N-bit bit vector that specifies which of the N virtual channels it is permitted to use throughout the physical network. That is, the vector contains a 1 in each bit position that corresponds to a permitted virtual channel and 0s in all other positions. The dispersion codes must be assigned so that for every pair of virtual networks, x and y, the bit vector corresponding to x contains a 1 in at least one position at which the bit vector corresponding to y contains a zero.

When a packet arrives at a node in a network that employs dispersion routing, the virtual channel to be used for the next hop of the packet's route is determined by intersecting two sets represented as bit vectors. A dispersion code vector A describes the set of virtual channels the packet is allowed to use (that is, it serves as a channel allocation vector), and a busy vector B describes the set of channels available.

As illustrated in Figure 15, each virtual network is associated with an N-bit dispersion code vector, A. Here N is the number of virtual channels multiplexed on each physical channel in the network. The dispersion code vector A associated with a particular virtual network, x, has C_x bits set to indicate the subset of virtual channels on which x is permitted to route. The dispersion code vectors for two virtual networks, x and y, have at most S bits in common, S being less than C_x and C_y .

The dispersion code vectors for all of the virtual networks may be stored in a table 102 as shown in Figure 15. If each virtual network is permitted to route on a relatively large fraction of the virtual channels, the table is most efficiently encoded directly as bit vectors as shown in the Figure. On the other hand, if each virtual network is restricted to a small fraction of the total channels, the vectors are more efficiently stored in compressed form as a list of channel indices. For example, the vector 0000100001001000 can be represented as the list of indices 3,6,11. Here, 3, 6 and 11 correspond to the positions of the 1 bits in the binary vector, the rightmost bit being in position 0. Representing the vector as indices requires 12 bits for three four-bit indices rather than 16 bits for the full vector. Alternatively, the table may be dispensed with entirely and the channel allocation vectors derived from the virtual network index (in this case the address of the destination node) using combinational logic. For example, if the virtual network is assigned one channel for its x

-31-

coordinate, one for its y coordinate, and one for its z coordinate, as discussed below, the coordinates of the destination address may be directly decoded to generate the dispersion code vector.

At any point in time, the state of the virtual channels associated with a physical channel is recorded in a set of vectors. A bit of the busy vector, B , is set if the corresponding virtual channel is currently assigned to carry a packet and, being busy, is not available to handle a new packet.

When a packet traveling on virtual network x arrives at a node, a routing decision is made to select an output port, p . Within the VC state table logic 80 (Figure 11B), corresponding to that output port, the complement of the busy vector of the output port $\overline{B_p}$ is then ANDed at 104 with the dispersion code vector A_x for virtual network x to generate a vector of candidate virtual channels, $V = \overline{B_p} \wedge A_x$. If V is the null vector (all zeros), no channel is available and the request is queued to be retried when a virtual channel for port p becomes available. If V is non-null, an arbiter 106 selects one of the non-zero bits of V , the corresponding channel is allocated to the packet, and the corresponding bit of B_p is set to mark the channel busy. Once the virtual channel is assigned, transport of flits takes place as described above.

Beyond avoiding the spread of congestion, which slows data transfers, to other virtual networks through shared virtual channels, care must be taken in assigning dispersion codes to avoid channel-dependent deadlocks which can halt transfers. Assignment of dispersion codes guaranteed to be deadlock-free, for 1-D and 2-D networks and then for a 3-D torus network follow.

Consider a 1-D bidirectional ring network as shown in Figure 13. In each direction around the loop, the span of a virtual network (VN) is the set of physical channels used by the virtual network. With minimal routing, a message initiated at node 2, for example, would not follow the route 2-1-6-5-4 because there is a minimal route 2-3-4, so the span of each VN covers half of the channels in the cycle. In Figure 13, for example, the span of the VN rooted at the shaded node 4 in the clockwise direction consists of the three bold channels. Its span in the other direction consists of the three fine channels that run in the opposite direction.

-32-

In networks having a radix, k , of 5 or more (Figure 13 having a radix of 6) and unrestricted assignment of a single virtual channel to each message, a dependent-cycle of three VNs with overlapping spans on the ring can cause deadlock. Consider, for example, three messages covering the respective routes 1-2-3-4, 3-4-5-6 and 5-6-1-2, all being initiated simultaneously. The worm hole of the first message would begin at node 1 and extend through node 2 but would be retained at the input buffer of node 3 since the message at node 3 would have already secured the shared virtual channel. The resultant delay of the first message until completion of the second message beginning at node 3 would of itself be acceptable. However, the message beginning at node 3 would not pass node 5 until completion of the third message, and the third message would be blocked by the initial message. Thus, further advancement of each message is prohibited by completion of another message in the cycle which is itself prohibited by the first message; that is, there is deadlock.

With dispersion routing, where each destination defines a virtual network (VN) that is permitted to use C virtual channels (VC) with a maximum overlap of S virtual channels ($0 < S < C$) between any pair of VNs, $3F$ VNs (where $F = \text{floor}(C/S)$) are required to generate a deadlocked configuration as a packet must block on F separate blocked VNs to deadlock.

A sufficient condition to avoid deadlock in a single dimension is for each VN to have at least one virtual channel (VC) that it shares only with VNs that overlap either entirely or not at all. That result can be obtained by assigning, for each coordinate of the loop, a VC used only by the VN rooted at that coordinate, requiring an assignment of 6 virtual channels to the network of Figure 13. However, with minimal routing, mirror nodes 1 and 4, for example, which are opposite to each other in the network, can share the same virtual channels since minimal routes to those two nodes would never share physical channels. Over any link within the loop, a message destined to one node would be traveling in an opposite direction to a message destined for the mirror node. Accordingly, the loop of Figure 13 requires a minimum of 3 virtual channels, each shared by only mirrored virtual networks, to avoid deadlock. Additional virtual channels may then be assigned to the loop for

shared or unshared use. With this approach, each VN is always able to make progress (within one dimension) through its unshared VC.

It is possible to avoid deadlock with a less restrictive assignment of VCs to VNs since it is only necessary to break the deadlock at one point in the cycle.

5 In a multidimension network it is possible to deadlock even if all dimensions are individually deadlock free. Consider the two-dimensional case which is easily extended to three dimensions. A deadlock can form if a packet making a NW turn blocks on a packet making a WS turn which in turn blocks on a packet making a SE turn, which in turn blocks on a packet making an EN turn, which in turn blocks on
10 the original packet. This forms a cycle (NW,WS,SE,EN). As shown by C.J. Glass and L.M. Ni, "The Turn Model for Adaptive Routing," Proceedings of the 19th International Symposium on Computer Architecture, May 1992, pp. 278-287, a cycle can be broken, thus avoiding deadlock, by eliminating any one of the turns in the cycle. Though complete elimination of the north edge, for example, of the cycle
15 would break the cycle, it is sufficient that only the turn EN be eliminated. A turn WN or straight NN can be allowed without risking deadlock.

Assume a two-dimensional array in which each of the X and Y dimensions is made deadlock free by assigning virtual channels as discussed above. For example, in an 8 x 8 array, $4 + 4 = 8$ virtual channels would be assigned to the appropriate
20 virtual networks. Assignment would be defined by the coordinate of the destination address (mod $k/2$), where k is the number of nodes in the respective dimension. The (mod $k/2$) factor accounts for mirror nodes which share VCs. If minimal routing is used, each VN is itself deadlock-free because, in each quadrant about the destination node, only two directions, and hence only two of eight possible turns, are used. This
25 is illustrated in Figure 14. In the region NE of the destination node, for example, packets only travel S and W and hence only SW and WS turns are allowed. This is one turn from the clockwise cycle and one turn from the counterclockwise cycle. If VNs share VCs, however, deadlock can occur as the turns missing from one VN may be present in other VNs sharing the same VC.

30 As discussed above, the cycles which result in deadlock can be eliminated by disallowing a turn in the cycle. To break both clockwise and counterclockwise cycles, a turn in each direction must be eliminated. Further, it must be assured that a

-34-

turn eliminated in one virtual channel is not replaced by the same turn in another virtual channel to complete the cycle. Since disallowance of both turns of a quadrant in a VN would prevent reaching the destination from that quadrant, the disallowed turns must be from different quadrants. Although restrictive, the destination can be reached from any quadrant if only a single turn of that quadrant is disallowed.

Accordingly, a sufficient method for preventing inter-dimension deadlock is to (1) make the dimensions individually deadlock free and (2) to require that (a) each VN disallow one of the four turns in each of the CW and CCW directions in different quadrants, and (b) each VN have at least one VC that is shared only with VNs disallowing the same turn. This is quite restrictive as it forces two of the four quadrants about the destination node to route in dimension order. For example, if the NE turn were eliminated, messages to the destination from the SW would have to be routed first to the E and then N since only the EN turn would remain available.

A strategy that permits more flexible routing, but is a more costly one in terms of VCs, is to associate two VNs with each destination node, each disallowing all turns for a single quadrant, the quadrants being different for the two VNs. For example, one VN for all quadrants but the NW would disallow the SE and ES turns and one for all quadrants but the SE would disallow the NW and WN turns. VNs from each class can then share VCs without restriction as long as they remain deadlock free in each dimension independently. Diagonal quadrants were chosen in the example for balance, but any two quadrants might be disallowed for the two VNs.

One workable method for assigning VCs in two dimensions is as follows:

1. Each destination is assigned two virtual networks, one disallowing SE and ES turns and one disallowing NW and WN turns.
2. Each VN is assigned a VC associated with the x-coordinate of the VNs destination $(\text{mod } k_x/2)$, where k_x is the number of nodes in the x-dimension. Assigning this VC guarantees non-overlap and hence single-dimension deadlock freedom in the x-dimension.

3. Each VN is assigned a VC associated with the y-coordinate of the destination (mod $k_y/2$). This guarantees single-dimension deadlock freedom in the y-dimension.
 4. Any additional VC pairs are assigned arbitrarily subject to the restriction that
5 no more than S VCs are shared between any two destinations.
 5. The routing tables are built so that nodes in the NW quadrant of a destination are restricted to the VN that disallows NW/WN and nodes in the SE quadrant are restricted to the other VN. Nodes in the NE and SW quadrants may use either VN.
- 10 As an example, for a 2-D network of 64-nodes (8x8) this assignment requires a minimum of 8VC pairs (16VCs total).
- The number of VCs which are available is dependent on both the buffer space which is available and on the number of flit buffers per VC in each node. Increasing the velocity of the system with increased flit buffers per node decreases
15 the available virtual channels. Beyond the minimum VCs required as identified in steps 2 and 3 above, any available virtual channels can be assigned to specific VNs, thus limiting S and the effect of congestion in one VN on another VN, or they can be shared to improve utilization of VCs across the network. For example, suppose two VCs are assigned to each VN in steps 2 and 3 above, and 20 channels remain
20 available for assignment after those steps. All 20 can be shared by all VNs for $C = 22$, $S = 21$, or each of the 20 may be assigned for exclusive use of a single VN for $C = 3$ and $S = 1$.

Increasing the sharing increases the value S. The ratio C/S is an indication of how many VNs can themselves be blocked due to congestion without affecting other
25 VNs. The larger that ratio, the lesser the effect of congestion in one VN on other VNs.

To extend this approach to three dimensions we need to exclude additional turns to avoid 3-D inter-dimension cycles. However, we can accomplish this with just two VNs per destination as above. One VN excludes the turns associated with

the NWU (North, West, Up) octant (SE,ES,SD,DS,ED,DE) while the other excludes the turns associated with the SED (South,East,Down) octant.

An example 1024-node network organized as $8 \times 8 \times 16$ needs a minimum of $4+4+8=16$ VC pairs (32 VCs) to assign one VC pair to each symmetric mirrored pair of planes in the network.

When a single destination receives an excessive amount of traffic, all VCs associated with its two VNs will quickly saturate and back up to the source. To first approximation it is as if these VCs were removed from the network, though, in fact, saturation is less likely to affect nodes further from the destination. With the channel assignment suggested above, where each destination node has two VNs with three VCs each (one per dimension), a VN to a node that shares exactly one coordinate with the saturated destination has four VCs remaining upon which to route. A VN to a node that shares two coordinates has two remaining VCs..

DEFLECTION ROUTING

Deflection routing is another method for making traffic destined for different fabric outputs substantially non-blocking. With deflection routing all of the packets are allowed to share virtual channels without restriction. When a packet blocks, however, rather than waiting for the required virtual channel to become available, it is misrouted or "deflected" to the packet memory of the present fabric router's line interface. It is reinjected into the fabric at a later time. Because a packet destined for fabric output A is never allowed to block, it cannot indefinitely delay a packet destined for fabric output B.

Deflection routing has several properties that make it less desirable than using virtual networks to achieve isolation between packets destined for different outputs. First, deflection routing provides no backpressure. When an output becomes congested, packets destined for that output are simply deflected and the fabric inputs sending packets to the congested output remain unaware of any problem. Second, while there is no blocking, there is significant interference between packets destined for different outputs. If an output, A, is congested, the links adjacent to A will be heavily utilized and a packet destined for output B that traverses one of these links will have a very high probability of being deflected.

Third, the use of deflection routing greatly increases the bandwidth requirements of the packet memory as this memory must have sufficient bandwidth to handle deflected packets and their reinjection in addition to their normal input and output. Finally, deflection routing is limited by the finite size of the packet memory on each
5 line interface. Under very high congestion, as often occurs in IP routers, the packet memory may be completely filled with deflected packets. When this occurs, packets must be dropped to avoid interference and possibly deadlock.

While this invention has been particularly shown and described with references to preferred embodiments thereof, it will be understood by those skilled
10 in the art that various changes in form and details may be made therein without departing from the spirit and scope of the invention as defined by the appended claims. Those skilled in the art will recognize or be able to ascertain using no more than routine experimentation, many equivalents to the specific embodiments of the invention described specifically herein. Such equivalents are intended to be
15 encompassed in the scope of the claims.

For example, the event-driven allocation logic described in connection with Figures 11A, 11B and 12 is suitable for use in an internet switch fabric router such as that shown in Figure 8. It should be understood that the event-driven allocation logic is also suitable for use in a multicomputer router. For example, with reference
20 to Figure 8, using a multicomputer interface as the line interface circuit 48 in combination with the event-driven allocation logic forms a multicomputer router for a multicomputer system such as that shown in Fig. 4.

Furthermore, it should be understood that the event-driven allocation logic is suitable for assigning input physical channels to output physical channels directly.
25 Preferably, a single copy of the allocation logic is used. The logic is activated by the occurrence of an event.

Additionally, it should be understood that portions of the state vectors for the virtual channel state table 80 (see Figure 12) have been described as including individual bits for indicating particular information such as busy or wait
30 information. Other structures can be used in place of such bits such as scalar state fields that encode the information.

In connection with the event-driven allocation logic described in Figures 11A, 11B and 12, it should be understood that each input physical channel is shared by multiple input virtual channels, and each output physical channel is shared by multiple output virtual channels. The allocation logic is suitable for providing a
5 single virtual channel for each physical channel. In such a case, each input physical channel is used by only one input virtual channel, and each output physical channel is used by only one output virtual channel. As such, the state table logic essentially generates assignments that associate input physical channels with output physical channels.

CLAIMS

What is claimed is:

1. An internet router for coupling to a plurality of internet links, the router receiving data packets from the internet links, analyzing header information
5 in the data packets to route the data packets and forwarding the data packets on the internet links, the improvement wherein the router comprises a fabric of fabric links joined by fabric routers, the number of fabric links to each fabric router being substantially less than the number of internet links served by the internet router, the fabric links and routers providing data
10 communication between internet links through one or more hops through the fabric.
2. An internet router as claimed in claim 1 wherein a line interface to each internet link analyzes the header information in data packets received from the internet link to identify an output internet link through an internet routing
15 protocol and to determine through a fabric routing protocol a routing path through the fabric to the identified output internet link.
3. An internet router as claimed in claim 2 wherein the link interface defines the routing path through the fabric by including, in a header, a link definition of each successive link in the routing path, each fabric router along the routing
20 path storing an associated link definition from the header for forwarding successive segments of the packet.
4. An internet router as claimed in claim 1 wherein, between hops on fabric links, segments of packets are stored in fabric routers at storage locations assigned to virtual channels corresponding to destination internet links.
- 25 5. An internet router as claimed in claim 1 wherein the number of internet links served by the internet router is at least an order of magnitude greater than the

-40-

number of fabric links to each fabric router, and wherein the number of virtual channels per fabric router is substantially greater than the number of links to the fabric router.

6. An internet router as claimed in claim 1 wherein each fabric router comprises a plurality of buffers for defining virtual channels which share fabric links, the virtual channels and links forming virtual networks between internet router inputs and outputs in which congestion in one virtual network is substantially nonblocking to packets flowing through other virtual networks.
7. An internet router as claimed in claim 1 wherein an arbitration is performed at each fabric router to assign a packet to a virtual channel for output from the fabric router and to assign a virtual channel to an output fabric link from the fabric router.
8. An internet router as claimed in claim 1 wherein each fabric router comprises a crossbar switch.
9. An internet router as claimed in claim 8 wherein each fabric router comprises input buffers for receiving data packets from respective fabric routers and an internet link.
10. An internet router as claimed in claim 9 wherein each fabric router comprises an output controller for forwarding data along an output fabric link upon receipt of an indication that an input buffer is available at the opposite end of the fabric link.
11. An internet router as claimed in claim 1 wherein the fabric is a direct network.
12. An internet router as claimed in claim 1 wherein the fabric is a three-dimensional torus network.

-41-

13. A method of routing data packets between internet links comprising:
analyzing header information in the data packets to route the data
packets to output internet links; and
routing the data packets through a multihop fabric network of fabric
5 routers to the output internet links.
14. A method as claimed in claim 13 further comprising analyzing the header
information in data packets at a line interface to identify an output internet
link through an internet routing protocol and to determine through a fabric
routing protocol a routing path through the fabric network to the identified
10 output internet link.
15. A method as claimed in claim 14 wherein the link interface defines the
routing path through the fabric by including, in a header, a link definition of
each successive link in the routing path, each fabric router along the routing
paths storing a link definition from the header for forwarding successive
15 segments of the packet.
16. A method as claimed in claim 13 wherein, between hops on fabric links,
segments are stored in fabric routers at storage locations assigned to virtual
channels corresponding to destination internet links.
17. A method as claimed in claim 13 wherein each fabric router comprises a
20 plurality of buffers for defining virtual channels which share fabric links, the
virtual channels and links forming virtual networks between internet router
inputs and outputs in which congestion in one virtual network is substantially
nonblocking to packets flowing through other virtual networks.
18. A method as claimed in claim 13 wherein an arbitration is performed at each
25 fabric router to assign a packet to a virtual channel for output from the fabric
and to assign a virtual channel to an output fabric link from the fabric router.

-42-

19. A method as claimed in claim 13 wherein the data packets are routed between fabric links through a crossbar switch.
20. A method as claimed in claim 13 wherein the data packets are routed between internet links through a three-dimensional torus direct network.
- 5 21. A network comprising a plurality of fabric routers connected by fabric links, each fabric router comprising a plurality of buffers for defining virtual channels which share fabric links, the virtual channels and links forming virtual networks between network inputs and outputs, where each virtual network is associated with a set of network outputs, in which congestion in
10 one virtual network is substantially nonblocking to packets flowing through other virtual networks.
22. A network as claimed in claim 21 wherein, between hops on fabric links, segments are stored in fabric routers at storage locations assigned to virtual channels corresponding to network outputs.
- 15 23. A router for routing data packets, comprising:
input physical channels for receiving at least portions of the data packets;
output physical channels;
data buffers, coupled with the input and output physical channels, for
20 storing the portions of the data packets; and
control circuitry, coupled with the input and output physical channels and the data buffers, for generating channel assignments in response to queued events, and outputting the portions of the data packets through the output physical channels according to the generated channel assignments.
- 25 24. The router of claim 23, wherein the control circuitry assigns virtual channels to the data packets in response to the queued events.

25. The router of claim 24, wherein the control circuitry is shared by multiple virtual channels and activated to handle a particular virtual channel in response to an event.
26. The router of claim 1, wherein each output controller further includes:
5 a state table that maintains states of output channels; and
 a transport circuit that queues transport requests when the state table is accessed in response to the queued events, and forwards data packets through the output physical channel according to the queued transport requests.
- 10 27. The router of claim 1, wherein each output controller receives credit events from a downstream router and, in response to the received credit events, queues a transport request to transport a portion of a data packet over the corresponding output physical channel.
28. The router of claim 23, wherein the control circuitry is adapted to generate
15 physical channel assignments in response to the queued events.
29. The router of claim 28, wherein the control circuitry is shared by multiple virtual channels and activated to handle a particular virtual channel in response to an event.
30. The router of claim 23, further comprising:
20 a multicomputer interface coupled with an input physical channel and an output physical channel such that the router forms a multicomputer router for a multicomputer system.
31. The router of claim 23, further comprising:
 a line interface coupled with an input physical channel and an output
25 physical channel such that the router forms an internet switch fabric router.

32. The router of claim 23, wherein the control circuitry includes:
a state table that associates the output physical channels with input channels.
33. A method for routing data packets from input physical channels to output
physical channels, the method comprising the steps of:
- 5 receiving at least portions of the data packets over the input physical
channels;
generating channel assignments in response to queued events; and
outputting the portions of the data packets through the output
physical channels according to the generated channel assignments.
- 10 34. The method of claim 33, wherein the step of generating channel assignments
includes the steps of:
generating virtual channel assignments that assign virtual channels to
the data packets; and
generating physical channel assignments that assign the output
15 physical channels to the virtual channels, each of the assignments being
generated in response to queued arrival and credit events, the portions of the
data packets being forwarded over the output physical channels according to
the generated virtual and physical channel assignments.
35. The method of claim 34, wherein the step of generating physical channel
20 assignments includes the step of:
queuing transport requests in a transport queue, the transport requests
being requests to transfer the portions of the data packets through the output
physical channels.
36. A router for routing data packets, comprising:
- 25 input physical channels for receiving at least portions of the data
packets;
output physical channels;

data buffers, coupled with the input and output physical channels, for storing the portions of the data packets; and

control circuitry, coupled to the input and output physical channels and the data buffers, which generates assignments of data packets to virtual channels which are multiplexed onto physical channels, a first data packet
5 having access to a first set of virtual channels of size C_1 and a second data packet having access to a second set of virtual channels of size C_2 , there being an intersection of the first and second sets of size S where $0 < S < C_1$, and $S < C_2$.

- 10 37. A router as claimed in claim 36 wherein the first packet routes on a first virtual network and the second packet routes on a second virtual network.
38. A router as claimed in claim 37 wherein all packets routed on the first virtual network share the first set of virtual channels and all packets routed on the second virtual network share the second set of virtual channels.
- 15 39. A router as claimed in claim 37 wherein each virtual network carries packets destined for a particular set of destination nodes.
40. A router as claimed in claim 39 wherein turns are disallowed in each virtual network to prevent interdimension deadlock.
- 20 41. A router as claimed in claim 36 wherein the control circuitry includes output controllers that correspond to the output physical channels, each output controller having a state table that records states of output virtual channels and identifies input virtual channels holding the portions of the data packet, and state table logic that accesses the state table to assign output virtual channels.
- 25 42. A router as claimed in claim 41 wherein the state table includes a virtual channel allocation vector for each virtual network.

43. A router as claimed in claim 42 wherein the state table further comprises a busy vector which indicates virtual channels which are busy in the output physical channel, and an output virtual channel is selected from a combination of the virtual channel allocation vector and the busy vector.
- 5 44. A router as claimed in claim 36 wherein substantially all of the data packets route over a plurality of virtual networks, each virtual network i using a subset of size C_i of the virtual channels, the subset overlapping the subset of at least one other virtual network and sharing no more than $S < C_i$ virtual channels within the other virtual network.
- 10 45. A router for routing data packets, comprising:
input physical channels for receiving at least portions of the data packets;
output physical channels;
data buffers, coupled with the input and output physical channels, for
15 storing the portions of the data packets; and
control circuitry, coupled to the input and output physical channels and the data buffers, which generates assignments of data packets to virtual channels which are multiplexed onto physical channels, data packets sharing access to multiple virtual channels to which other data packets have access, a
20 first data packet sharing some but not all virtual channels with a second data packet.
46. A network of routers for routing data packets comprising, at each router:
input physical channels for receiving at least portions of the data packets;
25 output physical channels;
data buffers, coupled with the input and output physical channels, for storing the portions of the data packets; and
control circuitry, coupled to the input and output physical channels and the data buffers, which generates assignments of data packets to virtual

channels which are multiplexed onto physical channels, a first data packet having access to a first set of virtual channels of size C_1 and a second data packet having access to a second set of virtual channels of size C_2 , there being an intersection of the first and second sets of size S where $0 < S < C_1$, and $S < C_2$.

47. A method of routing data packets comprising:
- storing at least portions of data packets;
 - generating assignments of data packets to virtual channels which are multiplexed onto physical channels, a first data packet having access to a first set of virtual channels of size C_1 and a second data packet having access to a second set of virtual channels of size C_2 , there being an intersection of the first and second sets of size S where $0 < S < C_1$, and $S < C_2$; and
 - forwarding the stored portions of data packets on an output physical channel.

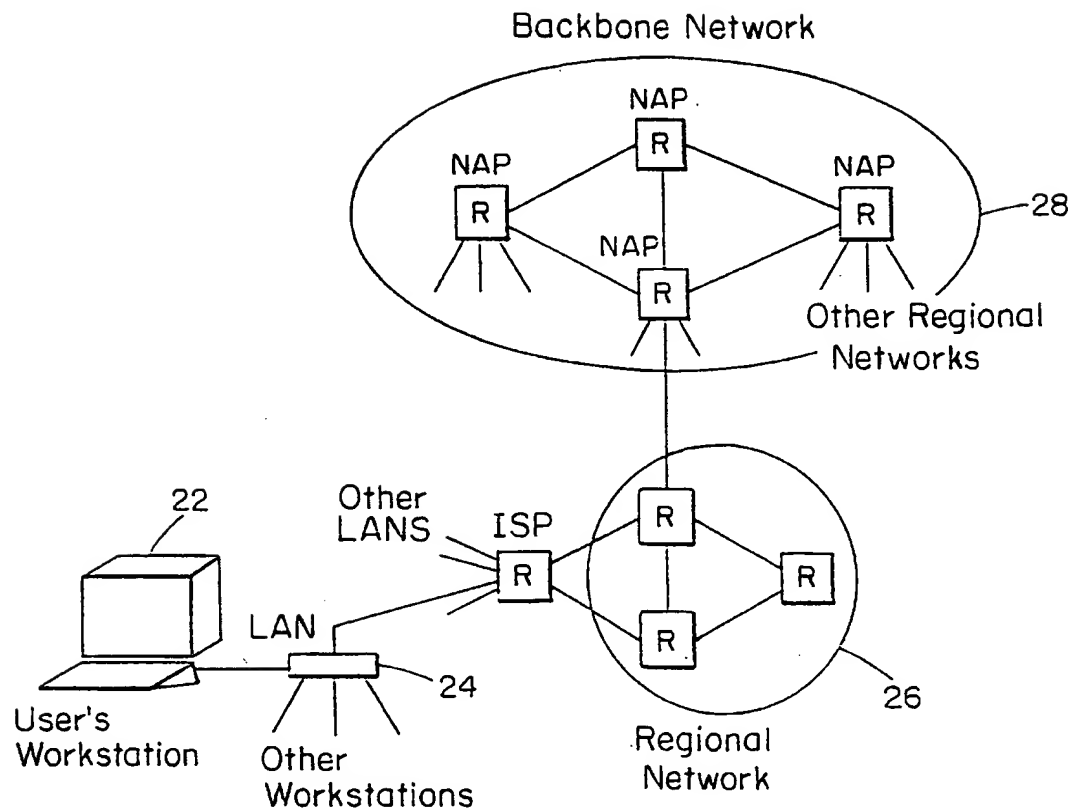


FIG. 1
PRIOR ART

2/11




FIG. 2
PRIOR ART




FIG. 3
PRIOR ART

3/11




FIG. 4




FIG. 5

4/11




FIG. 6

5/11




FIG. 7




FIG. 8




FIG. 9

7/11




FIG. 10A


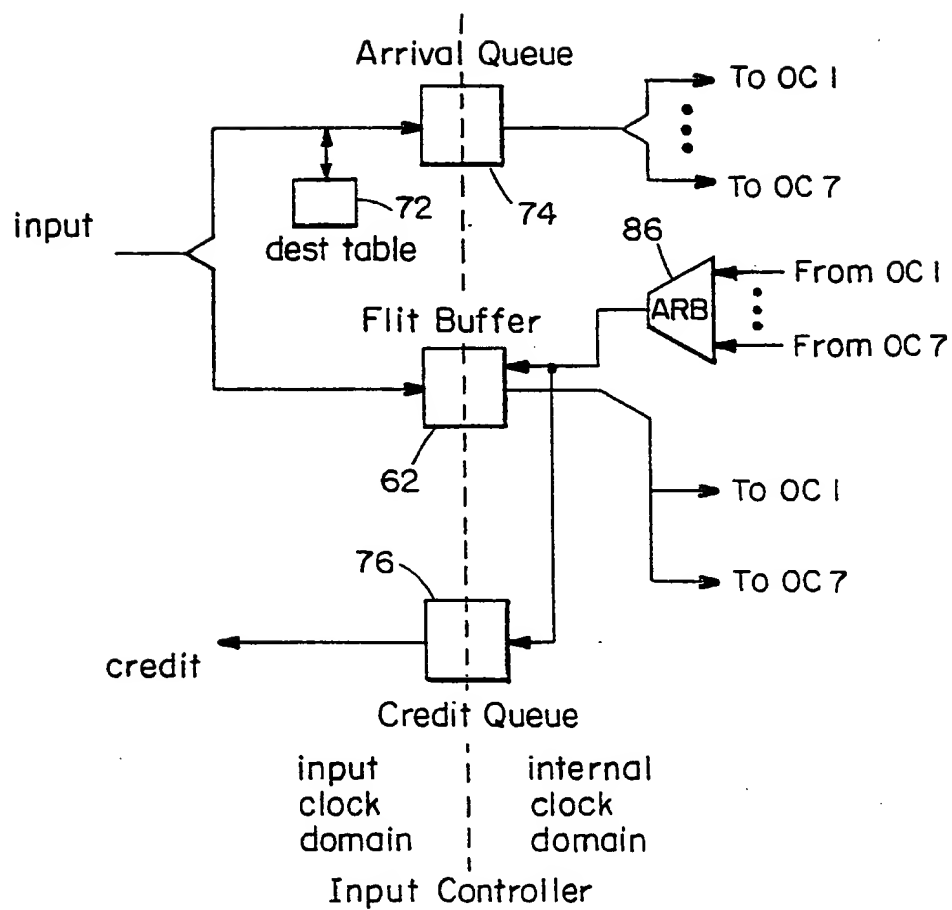
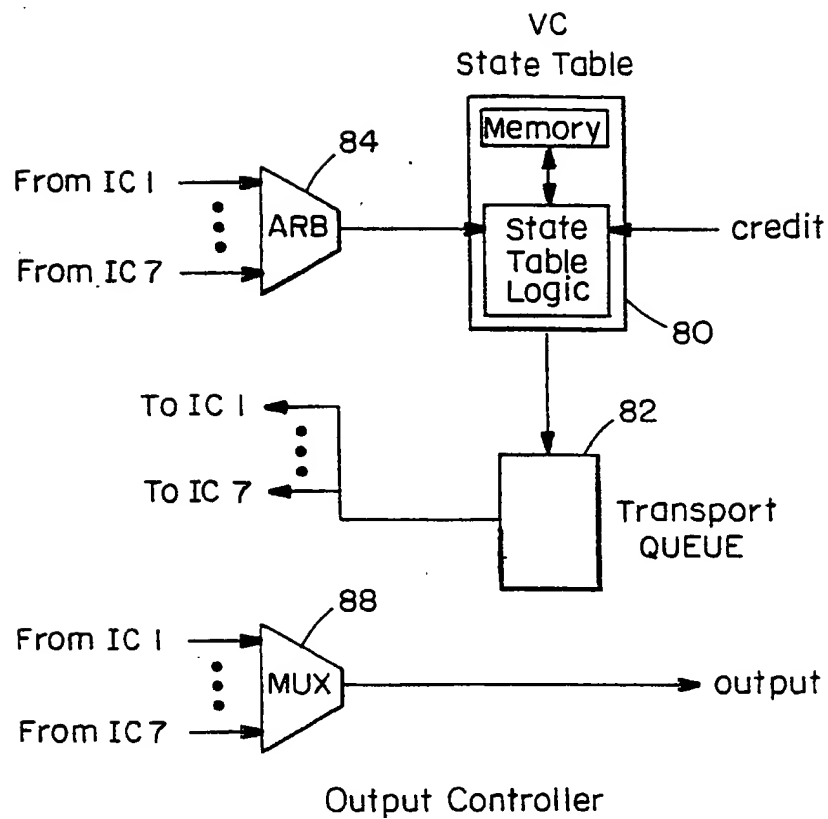


FIG. 10B

8/11

**FIG. IIA**

9/11

**FIG. 11B**

10/11

For Output K	Allocation Status (B = 0, 1, 2)	Input Controller (I = 1 to 7)	Waiting Input Controllers (W)	Credits (C)	Flits Present (P)
			7		
1					
2					
3					
•					
•					
•					
V					

state vector $s[V, K]$

FIG. 12




FIG. 13

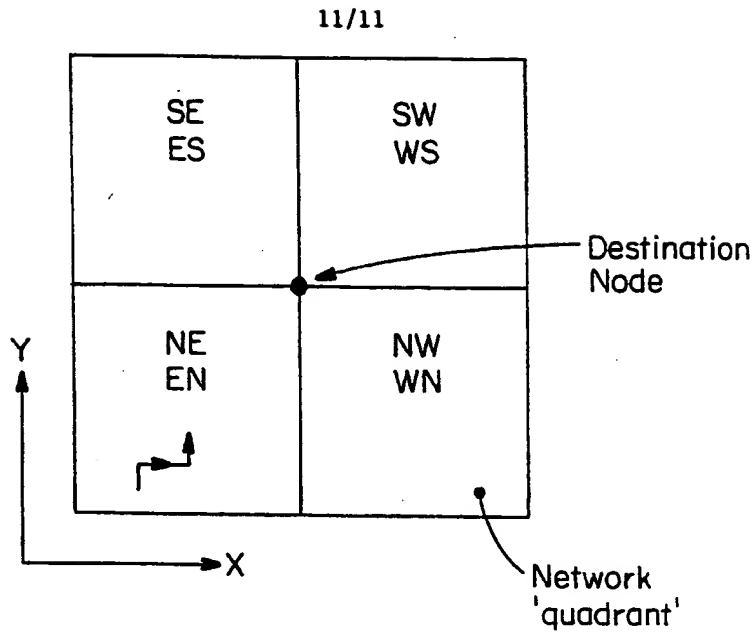


FIG. 14




FIG. 15

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 98/16762

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	PARULKAR G ET AL: "AITPM: A STRATEGY FOR INTEGRATING IP WITH ATM" COMPUTER COMMUNICATIONS REVIEW, vol. 25, no. 4, 1 October 1995, pages 49-58, XP000541650	1-4, 12-17, 20-22
A	see paragraph 2.1 see figure 2	8, 9, 19
Y	US 5 583 990 A (KESSLER RICHARD E ET AL) 10 December 1996	1-4, 12-17, 20-22
A	see column 2, line 8 - line 21 see column 10, line 61 - column 11, line 41 see column 14, line 13 - line 42 see column 18, line 4 - column 19, line 47	5, 21, 47

	-/--	



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

8 December 1998

Date of mailing of the international search report

18/12/1998

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Perez Perez, J

INTERNATIONAL SEARCH REPORT

Int tional Application No

PCT/US 98/16762

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	ROBERTAZZI T ET AL: "DEFLECTION STRATEGIES FOR THE MANHATTAN STREET NETWORK" COMMUNICATIONS - RISING TO THE HEIGHTS, DENVER, JUNE 23 - 26, 1991, vol. 3, no. __, 23 June 1991, pages 1652-1658, XP000277594 INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS see paragraph 2.3 ---	23-35
A	US 5 355 372 A (CHOWDHURY SHYAMAL ET AL) 11 October 1994 see abstract ---	23-46
A	YANG C S ET AL: "FAULT TOLERANT WORMHOLE ROUTING IN HYPERCUBE MULTICOMPUTERS" MICROPROCESSING AND MICROPROGRAMMING, vol. 35, no. 1 / 05, 1 September 1992, pages 667-672, XP000325181 see paragraph 2 -----	4-7, 16-18, 21,25, 29,36-47

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 98/16762

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5583990 A	10-12-1996	EP 0733237 A	25-09-1996
		JP 9506456 T	24-06-1997
		WO 9516240 A	15-06-1995
		US 5797035 A	18-08-1998
		US 5737628 A	07-04-1998
US 5355372 A	11-10-1994	JP 2546526 B	23-10-1996
		JP 7058754 A	03-03-1995
		JP 2546490 B	23-10-1996
		JP 6070350 A	11-03-1994